Efficient Dynamic Methods for Qualitative Constraint-based Spatial and Temporal Reasoning (Position Paper)

Michael Sioutis and Amy Loutfi

Centre for Applied Autonomous Sensor Systems (AASS)
Örebro universitet, Sweden
name.surname@oru.se

Abstract

We describe a research roadmap for going beyond the state of the art in Qualitative Spatial and Temporal Reasoning (QSTR). Our proposal involves researching novel local consistencies in the aforementioned discipline, and defining dynamic algorithms pertaining to these consistencies that can allow for efficient reasoning over streams of spatio-temporal information. Simply put, QSTR is a major field of study in Artificial Intelligence that abstracts from numerical quantities of space and time by using qualitative descriptions instead (e.g., precedes, contains, is left of); thus, it provides a concise framework that allows for rather inexpensive reasoning about entities located in space or time. Applications of QSTR are found in a plethora of areas and domains such as smart environments, planning, and unmanned aircraft systems.

1 Background

Qualitative Spatial and Temporal Reasoning (QSTR) is a major field of study in Artificial Intelligence, and in particular in Knowledge Representation & Reasoning, that deals with the fundamental cognitive concepts of space and time in an abstract, qualitative, manner. In a sense, this approach is in line with the qualitative abstractions of spatial and temporal aspects of the common-sense background knowledge on which the human perspective of physical reality is based. For instance, in natural language one uses expressions such as *inside*, *before*, and *north of* to spatially or temporally relate one object with another object or oneself, without resorting to providing quantitative information about these entities. More formally, QSTR restricts the vocabulary of rich mathematical theories that deal with spatial and temporal entities to simple qualitative constraint languages. Thus, QSTR provides a concise framework that allows for rather inexpensive reasoning about entities located in space and time and, hence, further boosts research and applications to a plethora of areas and domains that include, but are not limited to, ambient intelligence, dynamic GIS, cognitive robotics, spatio-temporal design, and qualitative model generation from video [31, 6, 14, 41].

As an illustration, the first constraint language to deal with time in a qualitative manner was proposed by Allen in [1], called Interval Algebra. Allen wanted to define a framework for reasoning about time in the context of natural language processing that would be reliable and efficient enough for reasoning about temporal information in a qualitative manner. In particular, Interval Algebra uses intervals on the timeline to represent entities corresponding to actions, events, or tasks. Interval Algebra has become one of the most well-known qualitative constraint languages, due to its use for representing and reasoning about temporal information in various applications. Specifically, typical applications of Interval Algebra involve planning and scheduling [3, 2, 25, 13], natural language processing [40], temporal databases [39, 8], multimedia databases [21], molecular biology [16] (e.g., arrangement of DNA segments/intervals along a linear chain involves particular temporal-like problems [5]), and workflow [23].

As another illustration, inspired by the success of Interval Algebra, Randell et al. developed the Region Connection Calculus (RCC) in [26], which studies the different relations that can be defined between regions in some topological space; these relations are based on the primitive relation of connection. For example, the relation *disconnected* between two regions x and y suggests that none of the points of region x connects with a point of region y, and vice versa. Two fragments of RCC, namely, RCC-8 and RCC-5 (a sublanguage of RCC-8 where no significance is attached to boundaries of regions), have been used in several real-life applications. In particular, Bouzy in [7] used RCC-8 in programming the Go game, Lattner et al. in [19] used RCC-5 to set up assistance systems in intelligent vehicles, Heintz et al. in [17] used RCC-8 in the domain of autonomous *unmanned aircraft systems* (UAS), and Randell et al. in [27] used a particular discrete domain counterpart of RCC-8 (called *discrete meterotopology*) to correct segmentation errors for images of hematoxylin and eosin (H&E)-stained human carcinoma cell line cultures. Other typical applications of RCC involve robot navigation, high level vision, and natural language processing [6].

2 Purpose and aims

Our proposal is to research novel local consistencies in the context of Qualitative Spatial and Temporal Reasoning (QSTR), with an emphasis on defining dynamic algorithms pertaining to these consistencies that can allow for efficient and flexible handling of incrementally available spatio-temporal information by means of real-time computing and just-in-time inference. Notably, the state of the art in QSTR lacks such dynamic algorithms (more to follow in Sections 3 and 4) and, hence, falls short of being practical for highly active applications [31].

To facilitate discussion, we first recall the notion of a *Qualitative Constraint Network* (QCN), which is a structure that is used to represent and reason about qualitative information, and is based on a finite set B of *jointly exhaustive and pairwise disjoint* binary relations defined over an infinite domain D (such as a topological space or the real line), called the set of *atoms* [20].

Definition 1. A *qualitative constraint network* (QCN) is a tuple (V, C) where:

- $V = \{v_1, \ldots, v_n\}$ is a non-empty finite set of variables, each representing an entity;
- and C is a mapping $C : V \times V \rightarrow 2^{\mathsf{B}}$ such that $C(v, v) = \{\mathsf{Id}\}$ for all $v \in V$, where Id denotes the *identity* atom, and $C(v, v') = (C(v', v))^{-1}$ for all $v, v' \in V$.

A subset of B (item of 2^{B}) denotes a relation encoding possible atoms, only one of which may hold between two entities. Hence, 2^{B} represents the total set of spatial or temporal relations. An example of a QCN is shown in Figure 1a; for simplicity, self-loops corresponding to relation $\{\mathsf{Id}\}$, and converse relations, are not depicted.

Given a QCN $\mathcal{N}$, we are particularly interested in its *satisfiability* problem, which is the problem of deciding if there exists a spatial or temporal interpretation of the variables of $\mathcal{N}$ that satisfies its constraints, such an interpretation being called a *solution* of $\mathcal{N}$ (an example of a solution

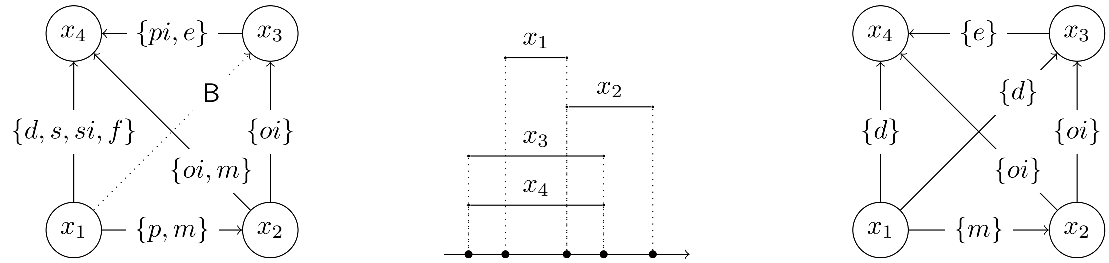

(a) A satisfiable QCN $\mathcal{N}$ (b) A solution σ of $\mathcal{N}$ (c) A scenario $\mathcal{S}$ of $\mathcal{N}$

Figure 1: Figurative examples of QCN terminology using Interval Algebra; symbols p, e, m, o, d, s, and f correspond to the atoms *precedes*, *equals*, *meets*, *overlaps*, *during*, *starts*, and *finishes* respectively, with $\cdot i$ denoting the converse of $\cdot$ (for a total of 13 atoms [1] as $ei = e$)

for a QCN of Interval Algebra is shown in Figure 1b). Other fundamental reasoning problems include the *minimal labelling* (or *deductive closure*) problem and the *redundancy* problem [29]. The minimal labelling problem is the problem of finding the strongest implied constraints of $\mathcal{N}$, and the redundancy problem is the problem of determining if a given constraint in $\mathcal{N}$ is entailed by the rest of $\mathcal{N}$ (that constraint being called redundant, as its removal does not change the solution set of the QCN). In general, for most qualitative constraint languages the satisfiability problem is NP-complete. Further, the redundancy problem, the minimal labelling problem, and the satisfiability problem are equivalent under polynomial Turing reductions [16].

Coming back to the discussion about the purpose of our plan of action, we want to push the envelope in QSTR by providing novel and efficient methods for tackling the aforementioned fundamental reasoning problems, but also—and most importantly—for tackling dynamic variations thereof, i.e., problems stated in terms of changing input data. To this end, we will pursue research in the following two directions:

- we will explore new local consistencies for QCNs and, in particular, local consistencies that rely upon singleton checks, i.e., constraint checks involving a temporary assignment of a singleton relation (defined by an atom) to the constraint at hand (for example, all the constraints in Figure 1c are defined by singleton relations), and that exploit the structure of the constraint graph of a given QCN, similarly to ongoing research in traditional constraint programming (which considers finite domains and, hence, different propagation techniques) [44, 24];

- we will define efficient algorithms pertaining to the new local consistencies, emphasizing on their dynamicity, which will enable them to be readily available for time-critical applications and dynamical systems that are subject to uncertainty and perturbation of the input data; furthermore, as existing state-of-the-art algorithms for enforcing current local consistencies are static in nature [34, 30, 35, 22, 32, 37], i.e., they operate on fixed input data (this will be stressed again in more detail in Section 4), we will provide dynamic variants of those algorithms as well and, consequently, fill this research void, which has occurred due to rapid advances in QSTR over the past few years missing out on that aspect.

In summary, our aim is to theoretically establish new techniques to efficiently handle qualitative spatial and temporal information in a dynamic setting, but also develop readily available tools for putting these techniques into practice.

3 State of the art

In the context of Qualitative Spatial and Temporal Reasoning, and with regard to local consistencies and algorithms for enforcing them in particular, the state of the art has been established by the main author and close collaborators in the published works of [34, 30, 35, 22, 32, 37]. We will briefly go over these key references in what follows.

In [34, 30] the authors propose a new local consistency (at that time) in the context of qualitative constraint-based reasoning that serves as the counterpart of *directional path consistency* in traditional constraint programming [12] or quantitative temporal reasoning [11], and is mainly distinguished by the fact that the involved consistency notions are tailored to handle infinite domains and qualitative relations. This local consistency is called *directional partial closure under weak composition* and is denoted by $\overleftarrow{\diamond}_G$*-consistency*. In particular, $\overleftarrow{\diamond}_G$-consistency entails consistency for all *ordered* triples of variables of a QCN that correspond to triangles of a given graph G. This ordering can be specified by a bijection between the set of the variables of a QCN and a set of integers, and can be chosen randomly or via an algorithm or heuristic.

Definition 2. A QCN $\mathcal{N} = (V, C)$ is $\overleftarrow{\diamond}_G$*-consistent* with respect to a graph $G = (V, E)$ and an ordering $(\alpha^{-1}(0),\alpha^{-1}(1),\ldots,\alpha^{-1}(n-1))$ defined by a bijection $\alpha : V \rightarrow \{0, 1, \ldots, n-1\}$ iff for all $v_i, v_j, v_k \in V$ such that $\{v_k, v_i\}, \{v_k, v_j\}, \{v_i, v_j\} \in E$ and $\alpha(v_i), \alpha(v_j) < \alpha(v_k)$ we have that $C(v_i, v_j) \subseteq C(v_i, v_k) \diamond C(v_k, v_j)$, where symbol $\diamond$ denotes the *weak composition* operation [20].

The authors then proceed to prove that $\overleftarrow{\diamond}_G$-consistency solves the satisfiability problem for a certain class of qualitative relations, called a *distributive* class of relations. This work is further extended in [35] to include more theoretical and practical results concerning other fundamental reasoning problems as well, such as the problem of scenario extraction from a satisfiable QCN (where a scenario is defined to be a satisfiable atomic QCN, see also Figure 1c).

Then, in [22] the authors show how $\overleftarrow{\diamond}_G$-consistency can be used to efficiently achieve $\diamond_G$-consistency for a given QCN that is defined over a distributive class of relations, which is a stronger local consistency with implications in the minimal labelling and the redundancy problems. Specifically, $\diamond_G$-consistency can be seen as $\overleftarrow{\diamond}_G$-consistency when the notion of ordered triples of variables is not taken into account, i.e., $\overleftarrow{\diamond}_G$-consistency entails consistency for *all* triples of variables of a QCN that correspond to triangles of a given graph G.

Definition 3. A QCN $\mathcal{N} = (V, C)$ is $\diamond_G$*-consistent* with respect to a graph $G = (V, E)$ iff $\forall \{v_i, v_j\}, \{v_i, v_k\}, \{v_k, v_j\} \in E$ we have that $C(v_i, v_j) \subseteq C(v_i, v_k) \diamond C(v_k, v_j)$.

In [32] the authors build upon the work of [22] and demonstrate how $\overleftarrow{\diamond}_G$-consistency can be used to efficiently achieve $\diamond_G$-consistency for *any* given QCN and not just for a QCN defined over a distributive class of relations. To this end, they exploit the notion of *abstraction* for QCNs, which is an idea adopted from concepts of abstract interpretation [9]. In particular, a QCN is typically abstracted by relaxing some of its constraints in order to satisfy some defined property.

Finally, in [37] the authors define a singleton-check-based local consistency that is *strictly stronger* than any of the local consistencies known to date, called *collective singleton-check-based consistency* and denoted by $\blacklozenge^{\cup}_G$-consistency.

Definition 4. A QCN $\mathcal{N} = (V, C)$ is $\blacklozenge^{\cup}_G$*-consistent* with respect to a graph $G = (V, E)$ iff $\forall \{v, v'\} \in E$, $\forall b \in C(v, v')$, and $\forall \{u, u'\} \in E$ we have that $\exists b' \in C(u, u')$ such that $b \in C'(v, v')$, where $(V, C') = \diamond_G(\mathcal{N}_{[u,u']/\{b'\}})$, and where $\diamond_G(\cdot)$ is the closure of $\cdot$ under $\diamond_G$-consistency.

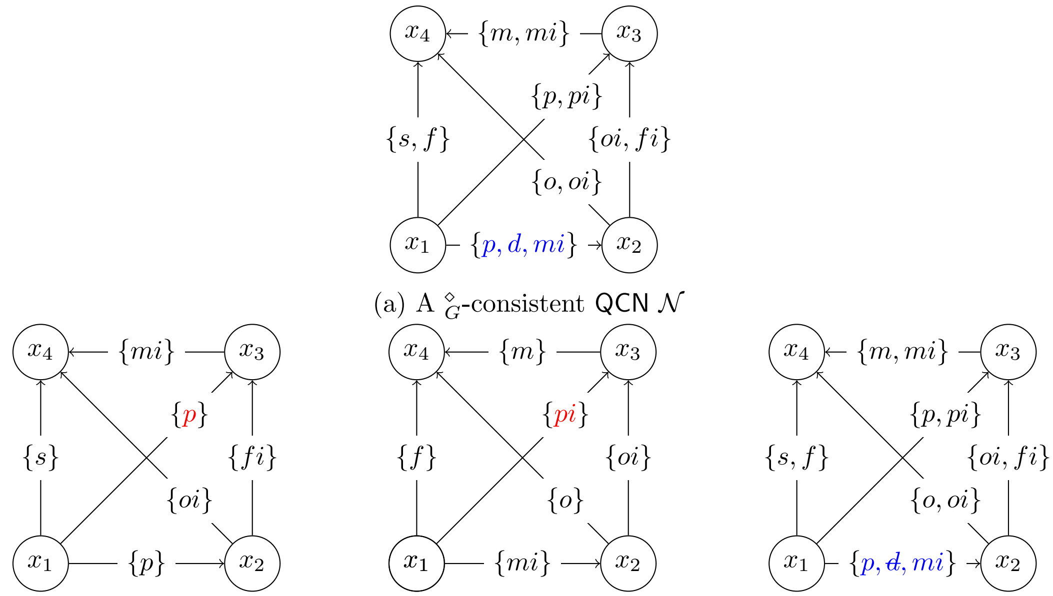

(a) A $\diamond_G$-consistent QCN $\mathcal{N}$

(b) $\mathcal{N}_1 = \diamond_G(\mathcal{N}_{[x_1,x_3]/\{p\}})$

(c) $\mathcal{N}_2 = \diamond_G(\mathcal{N}_{[x_1,x_3]/\{pi\}})$

(d) $\mathcal{N}_1 \cup \mathcal{N}_2$ $(\subset \mathcal{N})$

Figure 2: A $\diamond_G$-consistent QCN of Interval Algebra along with a demonstration of how enforcing $\blacklozenge^{\cup}_G$-consistency can further eliminate invalid atoms; G is the complete graph on $\{x_1, x_2, x_3, x_4\}$

A motivating example of the application of $\blacklozenge^{\cup}_G$-consistency is shown in Figure 2. As noted in [37] this local consistency can be essential for approximating satisfiability of QCNs and can play a crucial role in tackling the minimal labelling problem of a QCN in particular, as it is both strictly stronger and more efficient to enforce than the consistencies that had been utilized until that time to tackle the aforementioned problems; however, the exact performance gains that it may offer in the context of those problems have yet to be thoroughly experimentally studied.

4 Significance and scientific novelty

The significance and scientific novelty of the proposed research plan can already be drawn from the various applications of Qualitative Spatial and Temporal Reasoning (QSTR) detailed in Section 1 and the research directions described in Section 2. Specifically, by now the reader should be able to assert that QSTR is an active application area within Artificial Intelligence (AI), spanning several decades of research, and that fundamental scientific advances in that discipline are well adopted and appreciated by the research community and the industry. Nevertheless, in this section we delve into the particulars of our proposed research directions to pinpoint exactly how our proposal moves forward and innovates the current research frontier.

With regard to the direction concerning local consistencies that rely upon singleton checks, we would like to build upon the state-of-the-art local consistency of $\blacklozenge^{\cup}_G$-consistency, presented in Section 3, and define weaker variants of it, thus enriching the family of consistencies for QCNs. Specifically, the weaker variants will restrict singleton checks to the *neighbourhood* of the constraint in question. Early experiments in that direction have shown really promising results for constraint satisfaction problems (CSPs) in traditional constraint programming [44, 24], which is due to the fact that constraint revisions tend to propagate themselves to just neighbouring constraints. In that respect, we will also seek a balance between the strong theoretical properties that $\blacklozenge_G^{\cup}$-consistency offers, viz., that it is *strictly stronger* than any of the local consistencies known to date and can hence remove more unfeasible atoms in a given QCN than those consistencies (see [37, Section 4]), and the efficiency that should characterize algorithms for enforcing weaker variants of it. In particular, it will be interesting to investigate how good of an approximation certain variants of $\blacklozenge_G^{\cup}$-consistency can achieve in terms of pruning capability and consequent implications in the problems of satisfiability, minimal labelling, and redundancy.

Studying local consistencies by itself makes for a solid line of research, but, all things considered, in the end a local consistency is only as good as the algorithm that enforces it. This brings us to our second research direction, that of defining efficient algorithms pertaining to the new local consistencies with an emphasis on their dynamicity. As a first step, for the state-of-the-art algorithm that enforces $\blacklozenge_G^{\cup}$-consistency [37, Section 5], we would like to explore queuing strategies such that the singleton checks are applied in a more fruitful manner. In particular, it would make sense to prioritize certain singleton checks that are more likely to eliminate atoms anywhere in the network at hand, because this could unveil certain inconsistencies faster, but also lead to fewer constraint checks overall. Such strategies have been used to much success in the case of $\diamond_G$-consistency [43, 28]. These queuing strategies will be employed for the algorithms that will be designed to enforce the discussed weaker variants of $\blacklozenge_G^{\cup}$-consistency as well. Furthermore, dynamic algorithms will be developed to accommodate real-time computing and just-in-time inference for efficient and flexible handling of incrementally available spatial and temporal information For instance, let us consider the problem of *qualitative spatio-temporal stream reasoning*, i.e., the problem of incremental spatio-temporal reasoning over streams of information, studied in [17, 10]. This is an essential problem as, with the amount of data that is continuously produced, AI applications such as robotic systems are often tasked with reasoning about incrementally available information, and drawing relevant conclusions over such data flows and reacting to new situations with minimal delays is important. In both of those works, viz., [17, 10], the authors present approaches that rely upon the incremental functionality of the state-of-the-art algorithm at that time for enforcing $\diamond_G$-consistency,[1] which is described in [15, Section 3]. Although we have developed *much* more efficient algorithms for enforcing $\diamond_G$-consistency in a given QCN [22, 32] (see also the discussion in Section 3), such incremental functionality is not available in these algorithms for they are designed to operate on fixed input data; this is also the case for the state-of-the-art algorithms for enforcing the rest of the consistencies detailed in Section 3, viz., $\overleftarrow{\diamond}_G$-consistency and $\blacklozenge_G^{\cup}$-consistency [34, 30, 35, 37]. Thus, obtaining dynamic variants of those algorithms is a critical task that needs to be resolved.

5 Preliminary results

Towards our end goal of making the significant contributions specified in Section 2, we have already implemented a *lazy* algorithm for efficiently approximating $\blacklozenge_G^{\cup}$-consistency for a given QCN $\mathcal{N} = (V, C)$ with respect to a graph $G = (V, E)$ [36], a high-level description of which is provided in Algorithm 1. The algorithm is called ℓPSWC$^{\cup}$, which stands for *lazy* $\cup$*-collective partial singleton-check-based closure under weak composition*, and its laziness is due to the fact that it takes a non-exhaustive approach and performs a singleton check only for a constraint that has been revised and put into the queue due to a previous singleton check for some other constraint; this leads to

[1] The authors refer to $\diamond_G$-consistency as *algebraic closure* in their work.

Algorithm 1: ℓPSWC$^{\cup}(\mathcal{N}, G)$

in : A QCN $\mathcal{N} = (V, C)$, and a graph $G = (V, E)$.
out : A sub-QCN of $\mathcal{N}$.

begin
 $\mathcal{N} \leftarrow \diamond_G(\mathcal{N})$;
 $Q \leftarrow E$;
 while $Q \neq \emptyset$ **do**
  $\{v, v'\} \leftarrow Q.pop()$;
  $(V, C') \leftarrow \bot^V$;
  foreach $b \in C(v, v')$ **do**
   $(V, C') \leftarrow (V, C') \cup \diamond_G(\mathcal{N}_{[v,v']/\{b\}})$;
  $C(v, v') \leftarrow C'(v, v')$;
  if $(V, C') \subset \mathcal{N}$ **then**
   foreach $\{u, u'\} \in E \setminus \{v, v'\}$ **do**
    if $C'(u, u') \subset C(u, u')$ **then**
     $C(u, u') \leftarrow C'(u, u')$;
     $C(u', u) \leftarrow C'(u', u)$;
     $Q \leftarrow Q \cup \{u, u'\}$;
 return $\mathcal{N}$;

a *non-unique* closure being obtained in general for an input QCN. The aim of ℓPSWC$^{\cup}$ is to be significantly more efficient than an exhaustive algorithm for enforcing $\blacklozenge^{\cup}_G$-consistency and, at the same time, yield excellent approximations of $\blacklozenge^{\cup}_G$-consistency.

We evaluated the performance of algorithm ℓPWPC$^{\cup}$ against the state-of-the-art algorithm for enforcing partial $\blacklozenge^{\cup}_G$-consistency, viz., PWPC$^{\cup}$ [37], with random Interval Algebra QCNs of model $\mathsf{BA}(\mathsf{n}, \mathsf{m})$ [33]. This model can generate random QCNs of n variables with an average number $|\mathsf{B}|/2$ of atoms per non-universal constraint and by use of a *preferential attachment* [4] value m for the respective constraint graphs. We generated 100 Interval Algebra QCNs of $\mathsf{n} = 150$ variables for all values of m ranging from 2 to 5 with a step of 1, as the phase transition region for this model is observed for $\mathsf{m} \approx 3$ or 4 [30]. Regarding the graphs that were given as input to our algorithms, the *maximum cardinality search* algorithm [42] was used to obtain triangulations of the constraint graphs of our QCNs. The choice of such chordal graphs was reasonable given their extensive use in the recent literature, as reviewed in [38, 18].

For the evaluation we used a machine with an Intel Core i5-6200U processor, 8 GB of RAM, and the Ubuntu 16.04 OS. All algorithms were coded in Python and run using the PyPy interpreter[2] under version 5.1.2 (implementing Python 2.7.10). Only one CPU core was used.

The results of our experimental evaluation are detailed in Tables 1 and 2, where a fraction $\frac{x}{y}$ in Table 1 denotes that an approach required x seconds of CPU time and performed y constraint checks per atom removals on average per dataset of networks during its operation. In short, with respect to computational effort, Table 1 shows that ℓPSWC$^{\cup}$ had a significant advantage over PWPC$^{\cup}$ in all cases and, in particular, that ℓPSWC$^{\cup}$ was up to 3 times faster than PWPC$^{\cup}$ on average for the more difficult instances; and with respect to pruning capability, Table 2 shows that, in most cases, ℓPSWC$^{\cup}$ performed almost the same pruning in the labels of a given QCN as PWPC$^{\cup}$. We note

[2] `https://pypy.org/`

Table 1: Evaluation of the computational effort of algorithms $\mathsf{PSWC}^{\cup}$ [37] and $\ell\mathsf{PSWC}^{\cup}$ with structured random Interval Algebra networks of model $\mathsf{BA}(\mathsf{n} = 150, \mathsf{m})$ [33]

	min		μ		max		σ	
m	$\mathsf{PSWC}^{\cup}$	$\ell\mathsf{PSWC}^{\cup}$	$\mathsf{PSWC}^{\cup}$	$\ell\mathsf{PSWC}^{\cup}$	$\mathsf{PSWC}^{\cup}$	$\ell\mathsf{PSWC}^{\cup}$	$\mathsf{PSWC}^{\cup}$	$\ell\mathsf{PSWC}^{\cup}$
2	$\frac{0.00s}{1}$	$\frac{0.00s}{1}$	$\frac{0.55s}{5k}$	$\frac{0.11s}{741}$	$\frac{1.89s}{10k}$	$\frac{0.49s}{3k}$	$\frac{0.26s}{2k}$	$\frac{0.06s}{266}$
3	$\frac{0.00s}{1}$	$\frac{0.00s}{1}$	$\frac{9.64s}{16k}$	$\frac{2.32s}{4k}$	$\frac{52.56s}{58k}$	$\frac{18.74s}{23k}$	$\frac{7.89s}{8k}$	$\frac{3.02s}{4k}$
4	$\frac{0.00s}{1}$	$\frac{0.00s}{1}$	$\frac{81.35s}{44k}$	$\frac{31.03s}{17k}$	$\frac{874.06s}{348k}$	$\frac{589.13s}{238k}$	$\frac{143.24s}{74k}$	$\frac{74.64s}{34k}$
5	$\frac{0.00s}{1}$	$\frac{0.00s}{1}$	$\frac{0.77s}{22}$	$\frac{0.04s}{2}$	$\frac{24.75s}{569}$	$\frac{1.43s}{42}$	$\frac{3.31s}{91}$	$\frac{0.17s}{6}$

Table 2: Evaluation of the pruning capability of algorithm $\ell\mathsf{PSWC}^{\cup}$ compared to that of $\mathsf{PSWC}^{\cup}$ [37] with the Interval Algebra networks used in Table 1; a percentage $x\%$ denotes that $\ell\mathsf{PSWC}^{\cup}$ removed $x\%$ less atoms than $\mathsf{PSWC}^{\cup}$

m	min	μ	max	med
2	0%	0%	0%	0%
3	0%	0.22%	3.94%	0%
4	0%	1.45%	74.43%	0%
5	0%	0%	0%	0%

that the constraints were processed in random order.

6 Conclusion

In this position paper we provided a research roadmap of what we think the next steps should be to go beyond the state of the art in Qualitative Spatial and Temporal Reasoning (QSTR) and, hence, "pursue the Holy Grail" in that context. We made the case for researching novel local consistencies in the aforementioned discipline, and defining dynamic algorithms pertaining to these consistencies that can allow for efficient reasoning over streams of spatio-temporal information.

Acknowledgments. We would like to thank Prof. Jean-François Condotta for an interesting exchange of ideas on the topics of local consistencies and constraint programming in general.

References

[1] James F. Allen. Maintaining Knowledge about Temporal Intervals. *Commun. ACM*, 26:832–843, 1983.

[2] James F. Allen. Planning as Temporal Reasoning. In *KR*, 1991.

[3] James F. Allen and Johannes A. G. M. Koomen. Planning Using a Temporal World Model. In *IJCAI*, 1983.

[4] Albert-László Barabási and Réka Albert. Emergence of Scaling in Random Networks. *Science*, 286:509–512, 1999.

[5] Seymour Benzer. On the Topology of the Genetic Fine Structure. *Proc. Natl. Acad. Sci. USA*, 45:1607–1620, 1959.

[6] Mehul Bhatt, Hans Guesgen, Stefan Wölfl, and Shyamanta Hazarika. Qualitative Spatial and Temporal Reasoning: Emerging Applications, Trends, and Directions. *Spatial Cognition & Computation*, 11:1–14, 2011.

[7] Bruno Bouzy. Les concepts spatiaux dans la programmation du go. *Revue d'Intelligence Artificielle*, 15:143–172, 2001.

[8] Cindy Xinmin Chen and Carlo Zaniolo. Universal Temporal Data Languages. In *DDLP*, 1998.

[9] Patrick Cousot and Radhia Cousot. Abstract Interpretation Frameworks. *J. Log. Comput.*, 2:511–547, 1992.

[10] Daniel de Leng and Fredrik Heintz. Qualitative Spatio-Temporal Stream Reasoning with Unobservable Intertemporal Spatial Relations Using Landmarks. In *AAAI*, 2016.

[11] Rina Dechter, Itay Meiri, and Judea Pearl. Temporal Constraint Networks. *Artif. Intell.*, 49:61–95, 1991.

[12] Rina Dechter and Judea Pearl. Network-Based Heuristics for Constraint-Satisfaction Problems. *Artif. Intell.*, 34:1–38, 1987.

[13] Jürgen Dorn. Dependable Reactive Event-Oriented Planning. *Data Knowl. Eng.*, 16:27–49, 1995.

[14] Krishna Sandeep Reddy Dubba, Anthony G. Cohn, David C. Hogg, Mehul Bhatt, and Frank Dylla. Learning Relational Event Models from Video. *J. Artif. Intell. Res.*, 53:41–90, 2015.

[15] Alfonso Gerevini. Incremental qualitative temporal reasoning: Algorithms for the Point Algebra and the ORD-Horn class. *Artif. Intell.*, 166:37–80, 2005.

[16] Martin Charles Golumbic and Ron Shamir. Complexity and Algorithms for Reasoning about Time: A Graph-Theoretic Approach. *J. ACM*, 40:1108–1133, 1993.

[17] Fredrik Heintz and Daniel de Leng. Spatio-Temporal Stream Reasoning with Incomplete Spatial Information. In *ECAI*, 2014.

[18] Jinbo Huang, Jason Jingshi Li, and Jochen Renz. Decomposition and tractability in qualitative spatial and temporal reasoning. *Artif. Intell.*, 195:140–164, 2013.

[19] Andreas D. Lattner, Ingo J. Timm, Martin Lorenz, and Otthein Herzog. Knowledge-based risk assessment for intelligent vehicles. In *KIMAS*, 2005.

[20] Gérard Ligozat and Jochen Renz. What Is a Qualitative Calculus? A General Framework. In *PRICAI*, 2004.

[21] Thomas D. C. Little and Arif Ghafoor. Interval-Based Conceptual Models for Time-Dependent Multimedia Data. *IEEE Trans. Knowl. Data Eng.*, 5:551–563, 1993.

[22] Zhiguo Long, Michael Sioutis, and Sanjiang Li. Efficient Path Consistency Algorithm for Large Qualitative Constraint Networks. In *IJCAI*, 2016.

[23] Ruopeng Lu, Shazia Wasim Sadiq, Vineet Padmanabhan, and Guido Governatori. Using a temporal constraint network for business process execution. In *ADC*, 2006.

[24] Anastasia Paparrizou and Kostas Stergiou. On Neighborhood Singleton Consistencies. In *IJCAI*, 2017.

[25] Richard N. Pelavin and James F. Allen. A Model for Concurrent Actions Having Temporal Extent. In *AAAI*, 1987.

[26] David A. Randell, Zhan Cui, and Anthony Cohn. A Spatial Logic Based on Regions and Connection. In *KR*, 1992.

[27] David A. Randell, Antony Galton, Shereen Fouad, Hisham Mehanna, and Gabriel Landini. Mereotopological Correction of Segmentation Errors in Histological Imaging. *J. Imaging*, 3(4), 2017.

[28] Jochen Renz and Bernhard Nebel. Efficient Methods for Qualitative Spatial Reasoning. *J. Artif. Intell. Res.*, 15:289–318, 2001.

[29] Jochen Renz and Bernhard Nebel. Qualitative Spatial Reasoning Using Constraint Calculi. In *Handbook of Spatial Logics*, pages 161–215. 2007.

[30] Michael Sioutis. *Algorithmic Contributions to Qualitative Constraint-based Spatial and Temporal Reasoning*. PhD thesis, Université d'Artois, 2017.

[31] Michael Sioutis, Marjan Alirezaie, Jennifer Renoux, and Amy Loutfi. Towards a Synergy of Qualitative Spatio-Temporal Reasoning and Smart Environments for Assisting the Elderly at Home. In *IJCAI Workshop on Qualitative Reasoning*, 2017.

[32] Michael Sioutis and Jean-François Condotta. Efficiently Enforcing Path Consistency on Qualitative Constraint Networks by Use of Abstraction. In *IJCAI*, 2017.

[33] Michael Sioutis, Jean-François Condotta, and Manolis Koubarakis. An Efficient Approach for Tackling Large Real World Qualitative Spatial Networks. *Int. J. Artif. Intell. Tools*, 25:1–33, 2016.

[34] Michael Sioutis, Zhiguo Long, and Sanjiang Li. Efficiently Reasoning about Qualitative Constraints through Variable Elimination. In *SETN*, 2016.

[35] Michael Sioutis, Zhiguo Long, and Sanjiang Li. Leveraging Variable Elimination for Efficiently Reasoning about Qualitative Constraints. *Int. J. Artif. Intell. Tools*, 2018. In press.

[36] Michael Sioutis, Anastasia Paparrizou, and Jean-François Condotta. A Lazy Algorithm to Efficiently Approximate Singleton Path Consistency for Qualitative Constraint Networks. In *ICTAI*, 2017.

[37] Michael Sioutis, Anastasia Paparrizou, and Jean-François Condotta. Collective Singleton-Based Consistency for Qualitative Constraint Networks. In *TIME*, 2017.

[38] Michael Sioutis, Yakoub Salhi, and Jean-François Condotta. Studying the use and effect of graph decomposition in qualitative spatial and temporal reasoning. *Knowl. Eng. Rev.*, 32:e4, 2017.

[39] Richard T. Snodgrass. The Temporal Query Language TQuel. *ACM Trans. Database Syst.*, 12:247–298, 1987.

[40] Fei Song and Robin Cohen. The Interpretation of Temporal Relations in Narrative. In *IJCAI*, 1988.

[41] P. A. Story and Michael F. Worboys. A Design Support Environment for Spatio-Temporal Database Applications. In *COSIT*, 1995.

[42] Robert Endre Tarjan and Mihalis Yannakakis. Simple Linear-Time Algorithms to Test Chordality of Graphs, Test Acyclicity of Hypergraphs, and Selectively Reduce Acyclic Hypergraphs. *SIAM J. Comput.*, 13:566–579, 1984.

[43] Peter van Beek and Dennis W. Manchak. The design and experimental analysis of algorithms for temporal reasoning. *J. Artif. Intell. Res.*, 4:1–18, 1996.

[44] Richard J. Wallace. Neighbourhood SAC: Extensions and new algorithms. *AI Commun.*, 29:249–268, 2016.